Envision Toolbox Manual

Chapter 4

Editimage

Image Display and Manipulation

Chapter 4 - Editimage

Figure 1: The editimage application provides a variety of features for image display and manipulation. Here, the "kitten" image is displayed ($SAMPLEDATA/data/images/kitten.xv, shorthand "images:kitten").

A. Introduction

Editimage is an interactive image display and manipulation tool. The features provided by editimage include:

- Transparent support of a variety of input files
- Use of clip masks
- Capturing of images from the screen
- An input file history list
- For multiband images, control over which band is displayed
- Normalization of colormaps

- Control over which columns of the map data are displayed as red, green, and blue
- Application of functions to the map data values to produce new colormaps
- Interactive editing of the image file comment
- Application of autocolor procedures and colormap operations
- Regular and irregular region of interest extraction
- Creation of new data sets as defined by an interactively specified function
- Zooming on the displayed image
- Display of pixel values surrounding the location of the pointer in the image
- Display of map data values surrounding the location of the pointer in the image
- Colormap manipulation via a pseudo-coloring mechanism
- Pixel clipping and pixel thresholding on the image

The master form of editimage displays the input image in a workspace. Above the workspace are "License", "Help", and "Quit" buttons, from which one may access the license agreement, display online help pages, and quit the program, respectively. To the left of the image workspace are two subform buttons labeled "Files" and "Options".

Use the "Files" subform button to display the "Input and Output Files" subform. From this subform, you may input and output files.

A number of features are available from the "Options" subform, which contains seven panes: "Capture", "Image List", "Display", "Comment", "Colormaps", "Extract ROI", and "Function".

Beneath the "Options" subform button are five more subform buttons, each of which will display the relevant subform.

The "Zoom" subform lets you zoom in on the image.

The "Pixel Values" subform prints pixel values as the pointer is moved across the image.

The "Map Values" subform prints map data values as the pointer is moved across the image.

The "Pseudo Color" subform allows you to interactively modify the colormap of the image.

The "Threshold" subform supports pixel clipping and thresholding on the image.

B. Command Line Options

The syntax for executing editimage from the command line is as follows:

```
% editimage
  [-i]        (infile)  input image filename
  [-c]        (infile)  input clip mask filename
  [-cmap]     (infile)  input alternate colormap file
  [-use_cmap] (boolean) use colormap?
              (default = true)
  [-zoom]     (float)   initial zoom factor
              (zoom > 0.0, default = 2)
  [-update]   (float)   initial update time
              (update >= 0.0, default = 0)
  [-visual]   (cycle)   Read Only or Read Write Visual
                1  "Read Only",
                2  "Read Write"
                [default: 1 "Read Only"]
```

Note that the standardized Khoros arguments [-V], [-U], [-usage], [-P], [-gui], [-form], [-x], [-y], [-A], [-a], and [-ap] are also supported by editimage.

[-i {input image filename}]

This is the input image to be displayed. If desired, you may start editimage with no arguments, and enter the input image from the Files subform. The input image file may be of any data type. With multiband images, the first (0th) band is displayed by default; the "Image List" pane of the "Options" subform can be used to display other bands if desired.

[-c {clip mask filename}]

If desired, a bit image may be input and used as a clip mask. It must be a Khoros VIFF file of data type BIT; if the VIFF image passed in is NOT of type BIT, it will be automatically converted without prompting or warning (this is only done internally - the file itself will not be changed). ALSO: if the VIFF data file passed in is a multiband image, the first (0th) band will be taken as the clip mask.

When an image is used as a clip mask for the currently displayed image, the only part of the image that will appear normally is that part defined by the pixels in the clip mask image that have a value of (1). All other parts of the current image will appear in the background color of the image window (usually black). Images work best as clip masks when they are the same size as the input data. `$SAMPLEDATA/data/masks/` contains two images that make good clip masks. To get the clip mask to go away, NULL out the "Clip Mask" parameter, and hit <cr>.

[-cmap {alternate colormap filename}]

If desired, another image may be input, and its colormap will be used with the currently displayed image, instead of the currently displayed image's colormap.

[-use_cmap {t/f}]

While the default action of editimage is to use the colormap associated with the currently displayed image, no colormap is required to be used. If the colormap associated with the currently displayed image is not wanted, AND the [-cmap] argument is not used, specify "-use_cmap 0", or "-use_cmap false". In this case, the default colormap associated with the screen on which editimage is being run will be used. NOTE: when "Use Colormap" is specified as false, the [-cmap_image] option detailed above WILL NOT TAKE EFFECT.

[-zoom {initial zoom factor}]
If you intend to use the Zoom feature to zoom in on the currently displayed image, and you know the zoomfactor that you will wish to use (the default is 2.0), you may specify the zoomfactor here. When you bring up the "Zoom" subform, the zoomfactor will be already set to your desired value. It must be specified as a floating point value greater than zero.

[-update {update time}]
Suppose that while you are displaying an image with editimage, the Khoros VIFF file that is currently displayed is over-written. When the [-update] flag is specified as some number greater than zero (the default), this indicates to editimage that it should automatically register the fact that the input file has been changed, and update the currently displayed image accordingly. The value of the [-update] flag specifies (in seconds) how often the input file is to be checked for any potential modifications. It should be specified as a floating point value which is greater than zero. When the value of the [-update] flag is zero (the default value), editimage will NOT check the input file for modifications.

[-visual {1/2}]
This argument specifies whether a Read Only or a Read Write visual is to be used by editimage. See the Envision Manual, Chapter 1, Introduction, Section ?.? for details on Read Only and Read Write visuals.

C. The Files Subform (Input/Output)

C.1. Supported File Formats

Currently supported file formats for input and output include:

1. Khoros 2.0 VIFF (viff)
2. Khoros 1.0 xvimage (xvimage)
3. GIF (gif)
4. PNM (pnm)
5. Sun Raster (rast)

6. X Bitmap (xbm)

7. X Pixmap (xpm)

8. X Window Dump (xwd)

9. AVS (avs)

10. ARF (arf)

11. Ascii (ascii)

12. encapsulated PostScript (eps) {for output only}

For more detailed information on these file formats, please see Section ?.?

C.2. Input Image

Use the "Input Image" input file selection to input a new image file. Type the filename of the desired image in the parameter box of the "Input Image" input file selection; when the filename is correct, hit <cr>, and the new image will be displayed. Alternatively, the file browser may be used to input a new image; click on the "Input Image" button to display the browser..paragraph With multiband images, the 0th (first) band is displayed by default; the "Image List" pane of the "Options" subform can be used to select other bands to be displayed. Note that if the colormap of the currently displayed image has been modified, it will be over-written at this time by the colormap of the new input image.

C.3. Input Clip Mask

Use the "Input Image" input file selection to input a data object containing a clip mask that will be used to clip the input image. When a clip mask is used with the input image, the only part of the image that will appear normally is that part defined by the values in the clip mask that have a value of (1). All other parts of the image will appear in the background color of the image window.

If a clip mask currently being used is no longer desired, simply blank out the filename and hit <cr>. Inputting a clip mask of NULL indicates that the current clip mask is no longer to be used.

Note that only data objects of type BIT can be used as clip masks. If the data object to be used as a clip mask is not of type bit, it will be internally converted to type bit, without prompting or warning, before it is used. Note that this conversion will not affect the original file in any way.

C.4. Input Colormap

This feature allows you to use a different colormap with the image being displayed, other than the one that is included in the data object which describes the image. This feature is often useful if you have developed a data object containing only a colormap, and wish to install it on various images. Colormaps from other images can also be used with the input image.

Each time you provide a new filename for "Input Colormap" and hit the <cr> key, the new colormap is installed on the input image. To restore the original colormap, enter the same filename as that which is being used for "Input Image". Note that when an image is saved after colormap manipulations have been done, the *new* colormap is the one that will be saved with the image.

C.5. Output File Format

Before writing out the image using the "Output Image" or "Output Screendump" options, select the desired output file format from the list of file formats provided in the "Output File Format" display list. IMPORTANT NOTE: the output file format specified here will also apply to any ROI's extracted using the "Extract ROI" pane of the "Options" subform.

C.6. Output Image

The image currently being displayed may be written to an output file; the output image file will have any colormap modifications that have been made to the image during the editimage session. First, use the "Output File Format" list selection to specify which file format you prefer. Then, fill in the name for the output file in the text parameter box of the "Output Image" output file selection, and hit the <cr> key. A pop-up message will tell you that the file has been written out.

C.7. Output Colormap

The colormap of the image as it currently appears may be written by itself to an output file. Note that an output file of this type *cannot* be read back into editimage using the [-i] option, as it will have no image data. It *can*, however, be read back into editimage using the [-cmap] option, so that the colormap output will be used with the image data provided by another file.

C.8. Output Screendump

A screendump is a literal copy of whatever is displayed with editimage at any given time, and will also include the colormap that is currently installed on the image.

To get a screendump of the currently displayed image, provide the desired filename in the parameter box of the "Screendump Output" output file selection and hit <cr>.

D. The Options Subform

The "Options" subform provides access to seven subforms:

1. The "Capture" pane allows you to capture images from the screen and display them. If desired, captured images may also be written to an output file.

2. The "Image List" pane provides a convenient way to redisplay an image that was displayed earlier in the session. This subform also lets you specify which band of an image is to be displayed, when multiband images are used.

3. The "Display" pane gives you control over the colormap that is used with the image; the colormap may be normalized, and you can set which of the map columns is associated with the red, green, and blue "guns" of the display. You can define the red, green, and blue values displayed to be a function of the map values if desired.

4. The "Comment" pane allows you to edit the comment stored with the input image file.

5. The "Colormaps" pane offers a number of autocolor procedures and colormap operations that may be installed on the displayed image to achieve different visual effects.

6. The "Extract ROI" pane subform supports region of interest extraction. A number of ROI shapes are supported, including rectangle, polyline, circle, ellipse, line, and freehand. Extracted ROI's may be displayed as signals, images, or surfaces. The ROI Extraction policy allows you to control whether the ROI extracted should be the area inside the ROI shape, the area outside the ROI shape, or the values on the outline of the ROI shape.

7. The "Function" pane allows you to specify a function which will be used to define a data set to be displayed as an image.

D.1. The Capture Image Pane

The Capture Image pane of editimage provides a convenient way in which to create new images from windows displayed on the screen. Captured images are automatically displayed in the editimage workspace. Captured images may be written to an output file, if desired; the Khoros 2.0 VIFF file format will be used with output files.

The *source* of image capture may be a specific window, or the screen as a whole. You may capture the entire source, or you may interactively specify a portion of the source to capture using a rubberbanding method.

The image capture feature supports capture of both 24-bit and 8-bit images. The number of planes captured depends on the display from which the image is captured. For example, on a 24-bit display, capture of both

8-bit and 24-bit images is possible; the captured image will be 8-bit if the source of the captured image was 8-bit, or 24-bit if the source of the captured image was 24-bit. On an 8-bit display, an 8-bit image will be captured.

The image capture feature *cannot* handle capture of areas containing multiple or split visual displays, or areas containing multiple colormaps.

There are a number of different methods which may be used to capture images:

1. "Entire Screen"
2. "Screen Area"
3. "Entire Window, Interactive Specification"
4. "Entire Window, Specification by Name/ID"
5. "Window Area, Interactive Specification"
6. "Window Area, Specification by Name/ID"

IMPORTANT NOTE: when a window is used as the image source, if the window indicated for image capture is obscured by any other windows, the image captured will also include the obscuring window(s).

The image capture feature of editimage may be used to capture images using a number of methods. These methods are as follows.

Entire Screen

This method provides a convenient way to capture the entire screen.

First, select "Use screen as source". Second, select "Capture ALL of window or screen". Third, click on the "Capture" button.

This will cause the entire screen to be captured and displayed. Note that the entire screen can also be captured interactively using the "Area of Interest on Screen" method.

Screen Area

This method allows you to capture a specified portion of the screen, regardless of the number of windows which it may include.

Select "Use screen as source". Make sure that no other flags are selected. Then, click on the "Capture" button.

You will be prompted to rubber-band around the area of the screen to be captured. To rubberband the area of interest, first click on the upper left hand corner of the desired portion of the screen; the cursor will change to look like the upper left hand corner of a box. *Keeping the mouse button down*, move the mouse to the lower right hand corner of the desired portion of the screen; the active portion will be indicated by rubber banding. The rubber-banded portion of the screen will be displayed in the editimage workspace as soon as you release the mouse button.

Entire Window, Interactive Specification

This method lets you capture a single window in its entirety from the screen, interactively specifying the window that you want.

First, select "Use (interactively indicated) window as source". Second, select "Capture ALL of window or screen". Third, click on the "Capture" button.

When the cursor changes to the cross, click on the window to be captured. The window will be automatically displayed in the editimage workspace.

Entire Window, Specification by Name/ID

This method lets you capture a single window in its entirety from the screen, specifying the window that you want by its name or ID.

First, select "Name/ID of source window". Then, fill in the name or ID of the desired window. If the name or ID of a window is unknown, you may use the X11 utility program "xwininfo" to obtain it. Next, select "Capture ALL of window or screen". Third, click on the "Capture" button.

The window will be automatically displayed in the editimage workspace.

Window Area, Interactive Specification

This method lets you capture part of a single window, interactively specifying the window that you want, and then interactively specifying the area of interest within that window.

First, select "Use (interactively indicated) window as source". Make sure that no other flags are selected. Then, click on the "Capture" button.

First, you will be prompted to click on the window which is to be captured. After you click on the desired window, you will then be prompted to rubber-band the area of interest within that window. When rubber-banding, the cursor will be constrained to locations within the specified window, making it impossible to capture any area outside the window.

Window Area, Specification by Name/ID

This method lets you capture part of a single window, specifying the the window that you want by its name or window ID, and then interactively specifying the area of interest within that window. Execute:

First, select "Name/ID of source window". Then, fill in the name or ID of the desired window. If the name or ID of a window is unknown, you may use the X11 utility program "xwininfo" to obtain it. Make sure that no other flags are selected. Then, click on the "Capture" button.

You will be prompted to rubber-band the area of interest within the named window. When rubber-banding, the cursor will be constrained to locations within the specified window, making it impossible to capture any area outside the window.

Output of Captured Image

If desired, you may output the captured image to an output file. Simply fill in the desired output filename in the "Output filename" selection, and hit <cr>. This will have the effect of capturing the image and displaying it in the workspace, just as clicking on the "Capture" button does; however, in addition to being displayed in the editimage workspace, the captured image will also be written out to the specified file in Khoros 2.0 VIFF file format.

Bitmap Format

Getimage supports output of images in bitmap format. When the "Output image in bitmap format?" flag is selected, the captured image will be dithered down to a bit (binary) image before the resulting image is output.

Inclusion of Window Manager Decoration

By default, when a window is captured, the window manager decoration is not captured with the image. Select "Grab the window manager decoration?" to specify that the window manager decoration is to be included as part of the output image.

D.2. The Image List Pane

The "Image List" pane of the "Options" subform keeps a history of all the images that have been input to editimage during the session; thus, it provides as a convenient way to redisplay an image that was displayed earlier in the session. This subform also lets you specify which band of an image is to be displayed, when multiband images are used.

Image List

The list selection on the left of the pane will contain the name of each image that has been input to editimage during the session. Click on the name of the desired image, and it will be re-input automatically. Note that any colormap changes that were initiated on the image after it was input will be preserved, including pseudocoloring, thresholding, and installation of autocolor procedures and colormap operations.

Band List

The list selection on the right of the pane will contain one entry labeled "Image Band X" for each band of the currently displayed image. For single-band images, the list will only have one entry, "Image Band 0", and clicking on the entry will have no effect. However, when a multiband image is input, you may change the band currently being displayed by simply selecting the number of the desired band from the "Band List" selection.

D.3. The Display Pane

The "Display" pane of the "Options" subform allows you to control the appearance of the image by changing the way the red, green, and blue values of the map columns are displayed.

First, it is assumed that the image being displayed has a minimum of three columns in its map data. Thus, the capabilities offered here are not very useful with bit images. Grey scale images are considered to have three identical map columns, with incremental values ranging from 0 to 255; thus, while it does "make sense" to use the capabilities on this pane with grey scale images, operations will not make a great difference in the appearance of the image.

At a minimum, to take advantage of the capabilities on this pane, it is necessary to have an input images with a standard colormap, where each pixel value in the value data indexes into the three map columns, where the first map column specifies the red values of the pixels, the second map column specifies the green values of the pixels, and the third map column specifies the blue values of the pixels.

A much greater amount of functionality can be obtained with images that have specialized map data, for example, those images produced by clustering algorithms used as input to spectrum, where there are many map columns in the map data, and there is not a specific rule for which map columns will dictate the red, green, and blue values of the colormap displayed.

D.3.1. Colormap Normalization

Then, set the "Type Of Normalization" to "Global" or "Local". With local normalization, the normalization procedure is performed locally using the maximum and minimum values of each of the red, green, and blue map columns. For example, if the minimum and maximum values in the red column are 0 to 1, and the minimum and maximum values in the green column are 100 to 200, the red column will be normalized such that the value 0 stays 0, but the 1 value becomes 255, while the green column will be normalized such that the 100 value becomes 0, and the 200 value becomes 255. For global normalization, the normalization procedure is performed globally over all of the three columns displayed as red, green, and blue. Thus, with global normalization, if the minimum and maximum values over all the red, green, and blue columns are 1 to 200, all three map columns will be normalized such that 1 becomes 0 and 200 becomes 255.

Finally, use the "Normalization Method" toggle to specify the method of normalization you prefer. When normalization is to be performed, this attribute defines the algorithm which is used in normalization. Normalization may be a simple stretch over the values that may be displayed on your workstation, or any of three variations on a standard deviation normalization may be used. For a simple normalization of values to within displayable intensity bounds, use "0 <= values <= Max Colors". Alternatively, images can have contrast increased or reduced using on of the standard deviation normalizations. "-1*(std.dev) <= values <= 1*(std.dev" will normalize values to within 1 standard deviation, "-2*(std.dev) <= values <= 2*(std.dev)" will normalize values to within 2 standard deviations, and "-3*(std.dev) <= values <= 3*(std.dev)" will normalize values to within 3 standard deviations.

D.3.2. Changing the Map Columns Displayed as Red, Green, Blue

By default, an image with a colormap has the first map column defining the red values of the colormap, the second map column defining the green values of the colormap, and the third map column defining the blue values of the colormap. However, when the input image has more than three map columns present, it is useful to be able to change which of the map columns are providing the values for red, green, or blue.

Clicking on "Change Red Map Col", "Change Green Map Col", or "Change Blue Map Col" will pop up a list of the map columns stored in the input file. You may then choose the map column that is to provide the colormap values for red, green, or blue.

By default, the red, green, and blue map columns are set to the first, second, and third map columns, producing the results that you would normally expect for an image with a colormap.

D.3.3. Defining Functions of Map Columns To Define Red, Green, Blue

With images having colormaps made up of more than three columns, such as those produced by clustering algorithms, it is often informative to be able to apply a function to the values in those map columns in order to define the values that will be displayed as each of the red, green, and blue columns. For example, in an image having 6 map columns, you might define the red intensity of pixel 10 as (map column 2)[10]+ (map column 3)[10]/ (map column 4) [10]. In this case, the "RED", "GREEN", and "BLUE" string selections may be used to define the function to be applied in order to produce the values that will be used as red, green, and blue. Functions must have *only* variable *M*, where *M* stands for "map column"; following each *M* must be a number *starting at zero* and ranging to *N-1*, where *N* is the number of map columns available in the image. For example, a valid function for red might be:
(M2 - M3)/(M2 + M3)
or
(M0 + M1 + M2)/(M3 - M4)
Equation evaluation follows the standard rules of precedence; evaluation proceeds from left to right, and use of parentheses is fully supported.

In a more simple use, equivalent to using the "Change Red Map Col", "Change Green Map Col", and "Change Blue Map Col" buttons, the red, green or blue functions may also be set to represent the map column which is to be used to specify the red, green, or blue values of the pixels in the image. For example, if you wanted the fourth map column to define the red values of the pixels in the image, (remember that map column numbering begins at 0), you could set the "Red" function to:
M3
You can also set the "RED", "GREEN", and "BLUE" functions to a constant, if desired. For example, if you wanted the red values of all the pixels in the image to be 200, you could enter "200" in the "RED" function.

D.4. The Comment Pane

The "Comment" pane of the "Options" subform allows you to create or modify the comment that is stored in the header of the image that is currently displayed; this feature is useful for documenting images. Khoros 2.0

VIFF images do not have a limit on the number of characters can be used in the comment.

Enter the desired comment in parameter box of the "Edit Comment" string selection. When the comment appears as desired, click on the "Change Comment" action button.

IMPORTANT: the image file must be output for the comment to be saved! Use the "Files" subform to output the image with the new comment.

D.5. The Colormaps Pane

The "Colormaps" pane of the "Options" subform offers a number of autocolor procedures that may be used to define a new colormap for the image, and a number of colormap operations that may be used to modify the current colormap of the image.

The different colormap operations and autocolor procedures available are detailed in Chapter 1, "Introduction", of the Envision manual.

D.6. The ROI Pane

The ROI pane of editimage provides region of interest extraction capabilities. To initiate the ROI extraction process, enter an output filename and hit <CR>, or click on the "Extract" button.

Output File
Provide the name of the output file to which the ROI is to be written. Note that the ROI does not *have* to be written to an output file; it may simply be displayed in a separate window (the ROI will be displayed in a separate window when "Display Extracted ROI" is set to "True"). Leave the filename as NULL if you do not want the ROI written to an output file.

ROI Extraction Policy
This list selection allows you to specify whether regions of interest will be defined by the area inside the shape, outside the shape, or by the outline of the shape.

Display Extracted ROI
Set this logical selection to TRUE if you want the extracted region of interest to be displayed in a separate window; change it to FALSE if you do not want to see the extracted ROI.

Extract Multiple ROI's
If desired, you may extract multiple regions of interest from the same image, and save them as a single data object.

ROI Presentation
The ROI Presentation dictates how the data object representing the extracted ROI is to be constructed. It may be output as a signal (2D plot), an image, or a surface (3D plot).

ROI Shape
This toggle is used to set the desired shape of the ROI. Supported ROI shapes include rectangle, line, polyline, circle, ellipse, and freehand.

Extraction Mode
The Extraction mode may be set to "Mouse" or "Keyboard". "Mouse" implies interactive extraction using the mouse, while "Keyboard" implies non-interactive extraction using the provided floating point selections to specify location & size coordinates.

D.6.1. ROI Shapes

The shapes supported for region of interest extraction include rectangle, line, polyline, circle, ellipse, and freehand. An ROI is defined both by the shape of the ROI, and by the ROI Extraction Policy. The ROI Extraction Policy may be one of "Inside Shape", "Shape Outline", or "Outside Shape". The following section summarizes the results of using the various ROI shapes with each of the ROI Extraction Policies.

Rectangle
The region of interest will be defined by a simple rectangle.

1. Inside Shape
The subimage defined by the rectangle is extracted.

2. Shape Outline
The bounding box of the rectangle is calculated, and all pixel values inside the bounding box but not on the outline of the rectangle are masked out.

3. Outside Shape
The pixel values within the rectangle are masked out, while the remainder of the image is preserved.

To extract a rectangular region of interest, rubberband the rectangle surrounding the region of interest in the input image. To rubberband the rectangle, click the mouse button on a corner of the desired region of interest. Holding the mouse button down, move the mouse to the opposite corner of the desired rectangle. The rectangle will refreshed as the mouse moves, or "rubberbanded" after the mouse. Release the mouse button when the rectangle appears as desired.

Polyline
The region of interest will be defined by a polyline. The first and last points of the polyline are joined in order to close the shape.

1. Inside Shape
The bounding box is calculated from the vertices of the polyline and all pixel values within the bounding box but not inside the polygon are masked out.

2. Shape Outline

The bounding box is calculated from the vertices of the polyline and all pixel values not on the polyline itself are masked out.

3. Outside Shape

All pixels within the polyline are masked out, while the rest of the image is preserved.

To extract a polyline region of interest, begin by clicking the mouse on one vertex of the desired polyline. As you move the mouse, a line will be rubberbanded after it. Continue to button click on each successive vertex, until the polyline appears as desired; double click to end the polyline. The first and last vertices of the polyline will be joined in order to define a closed region of interest.

Circle

The region of interest will be defined by a circle.

1. Inside Shape

The bounding box is calculated from the dimensions of the circle and all pixel values within the bounding box but not inside the circle are masked out.

2. Shape Outline

The bounding box is calculated from the vertices of the polyline and all pixel values not on the circle itself are masked out.

3. Outside Shape

All pixels within the circle are masked out, while the rest of the image is preserved.

To extract a circular region of interest, rubberband the circle defining the region of interest in the input image. To rubberband the circle, click the mouse button on the center of the desired region of interest. Holding the mouse button down, move the mouse until the circle has the desired radius; release the mouse button to finalize the circle.

Ellipse

The ellipse region of interest works similarly to that of a circle (see above).

To extract an ellipsoid region of interest, rubberband the ellipse defining the region of interest in the input image. Rubberbanding an ellipse is similar to rubberbanding a circle, except that the width and height of the ellipse will be changed with horizontal and vertical mouse movement.

Line

The region of interest will be defined by a line.

1. Inside Shape

The bounding box is calculated from the endpoints of the line; all pixel values within the bounding box but not on the line are masked out.

2. Shape Outline
Same as "Inside Shape" (see above).

3. Outside Shape
All pixels on the line itself are masked out, while the rest of the image is preserved.

To extract a line region of interest, rubberband the line defining the region of interest in the input image. To rubberband the line, click the mouse button on one endpoint of the desired region of interest. Holding the mouse button down, move the mouse to the other endpoint of the line. Release the mouse button when the line is of the desired length and direction.

Freehand
The region of interest will be defined by a freehand drawing. The begin and end points of the freehand drawing are joined in order to close the shape.

1. Inside Shape
The bounding box is calculated from the boundary of the drawing and all pixel values within the bounding box but not inside the freehand shape are masked out.

2. Shape Outline
The bounding box is calculated from the outline of the freehand drawing and all pixel values not on the outline itself are masked out.

3. Outside Shape
All pixels within the freehand drawing are masked out, while the rest of the image is preserved.

To extract a freehand region of interest, begin by clicking the mouse on an endpoint of the shape to be drawn. Holding the button down, move the mouse in the desired shape; a line will be drawn according to the movement of the mouse. Release the mouse button to end the freehand drawing. The beginning and ending point of the drawing will be joined in order to define a closed region of interest.

D.6.2. The ROI Extraction Policy

The ROI Extraction Policy can be set to "Inside Shape", "Shape Outline", or "Outside Shape"; this indicates whether regions of interest will be defined by the area inside the shape, outside the shape, or by the outline of the shape.

Masking is used to indicate which portions of the image are within the region of interest. Masked-out values will appear in black.

Since single or multiple ROI's can be extracted from the same image, the algorithm used for determining the mask value depends both on the ROI Extraction Policy, and on whether single or multiple regions of interest are being extracted from the image.

Inside/Single
When the ROI Extraction policy is set to "Inside Shape" and a single region of interest is extracted from the image, the mask values within the ROI are set to 1, while the mask values outside the ROI are set to 0.

Inside/Multiple
When the ROI Extraction policy is set to "Inside Shape" and multiple regions of interest are extracted from the image, the mask values within a ROI is set to the number indicating the order in which the ROI was created. For example, the mask values within the first ROI created are 1, the mask values within the second ROI created are 2, etc. The mask values outside the ROI are set to 0.

Outline/Single
When the ROI Extraction policy is set to "Shape Outline" and a single region of interest is extracted from the image, the mask values on the outline of the ROI shape are set to 1, while the mask values not on the ROI outline are set to 0.

Outline/Multiple
When the ROI Extraction policy is set to "Shape Outline" and multiple regions of interest are extracted from the image, the mask values on the outline of a ROI is set to the number indicating the order in which the ROI was created. For example, the mask values on the outline of the first ROI are 1, the mask values on the outline of the second ROI are 2, etc. The mask values not on the outline of any ROIs are set to 0.

Outside/Single
When the ROI Extraction policy is set to "Outside Shape" and a single region of interest is extracted from the image, the mask values within the shape are set to 0, while the mask values outside the ROI are set to 1.

Outside/Multiple
When the ROI Extraction policy is set to "Outside Shape" and multiple regions of interest are extracted from the image, the mask values within the all the shapes drawn are set to 1. The mask values outside the shapes drawn are set to 0.

D.6.3. Multiple ROI Extraction

If desired, you may extract multiple regions of interest from the same image, and save them as a single data object. Set "Extract Multiple ROIs" to "True" to enable this capability. Then, begin the ROI extraction process normally. After a region of interest is specified, the shape drawn on the image to delineate the ROI will not be erased as it is when single ROI's are extracted. Instead, the shape will remain on the image, and the next ROI may be extracted. When all ROI's have been specified as desired, click the *right mouse button* or *double-click* to indicate that you are done. The resulting data object will have the multiple ROI's specified combined.

Note that extractor will allow you to change the ROI Shape during multiple ROI extraction, so that ROI's of different shapes can be extracted into the resulting image. Editimage, on the other hand, does not allow the ROI Shape to be changed.

D.6.4. The ROI Presentation

The extracted ROI should be output according to the way you intend to visualize it. It may be output as a signal (2D plot), an image, or a surface (3D plot).

Signal

When the ROI Presentation is set to signal, the value data defined by the *outline* of the ROI shape will be extracted and saved, regardless of the ROI Extraction Policy. Only the data on the outline of the ROI shape is output; no masking is used. The output file is expected to be visualized as a 2D plot. When the shape is a rectangle, the data on the outline is extracted beginning at the upper left hand corner, and the outline is traversed in a clockwise direction. When the shape is a polyline, the data on the outline is extracted beginning at the first vertex of the polyline; the extraction process then proceeds through the rest of the vertices and ends at the last vertex of the polyline. When the shape is a circle or an ellipse, the extraction process begins at the exact top of the circle or ellipse, and proceeds in a clockwise direction around the shape. When the shape is a line, the extraction process begins at the first point of the line and proceeds to the end of the line.

Image

When the ROI Presentation is set to image, the value data defined by the *inside, outline, or outside* of the ROI shape will be extracted and saved, depending on the ROI Extraction Policy. Regions excluded by the ROI will be masked out. See "ROI Shape" for more details. The output file is expected to be visualized as an image.

Surface

When the ROI Presentation is set to surface and the ROI Extraction Policy is set to "Inside Shape", the value data defined by the scanlines inside the ROI shape will be extracted and displayed as a 3D plot; values outside the ROI shape to the bounding box (if applicable), are set to the minimum value, or "floored". When the ROI Extraction Policy is set to "Shape Outline", the value data defined by the ROI shape will be extracted and displayed as a 3D plot; values inside and outside are "floored". When the ROI Extraction Policy is set to "Outside Shape", values inside the ROI shape are "floored". Note that depending on the size of the image, a 3D plot of the "Outside Shape" may be indistinguishable from a 3D plot of the entire image.

D.6.5. Extraction Mode

When set to "Mouse", ROI extraction is interactive; the mouse is used inside the image to do the extraction. When set to "Keyboard", ROI extraction is non-interactive; the appropriate floating point selections for location & size specification of the current ROI shape will be mapped at the bottom of the subform, and the "Extract" button used to extract the desired shape. For each ROI shape, the interactive or "Mouse" mode of extraction has already been described in the preceding section entitled "ROI Shapes" so that information will not be repeated here. The following is a description of noninteractive, or "Keyboard" extraction.

Rectangle
With keyboard extraction of a rectangle, the floating point selections "Upper Left X", "Upper Left Y", "Width", and "Height" will be mapped. Use these parameters to specify the upper left corner of the rectangle, as well as its width and height; click on "Extract" to extract the rectangle.

Polyline
Polylines cannot be extracted using the keyboard method.

Circle
For keyboard extraction of a circle, the floating point selections "Center X", "Center Y", and "Radius" will be mapped. Use these parameters to specify the center and radius of the circle; click on "Extract" to extract the circle.

Ellipse
For keyboard extraction of a circle, the floating point selections "Center X", "Center Y", and "A" and "B" will be mapped. Use these parameters to specify the center, a and b values of the ellipse; click on "Extract" to extract the ellipse.

Line
For keyboard extraction of a line, the floating point selections "Begin X", "Begin Y", and "End X" and "End Y" will be mapped. Use these parameters to specify the two end points of the line; click on "Extract" to extract the line.

Freehand
Freehand ROI shapes cannot be extracted using the keyboard method.

E. Zooming

The Zoom subform allows you to zoom in on an entire image or on a selected region of the image.

About the Zoom Window
The zoom window, located at the bottom of the Zoom subform, will display the portion of the displayed image over which the pointer is centered. The zooming procedure is performed via pixel replication.

Zoom Factor
Enter the desired zoom factor in the parameter box of the "Zoom Factor:" float selection. Hitting the <cr> key will update the zoom factor, and the image in the zoom window will now be magnified by the new zoom factor. Warning: When the Interpolation type is set to pixel replication, zooming by a factor of less than one will may ambiguous results.

Zoom Window Cursor
The zoom window, by default, has a "dummy" cursor displayed that reflects the position of the

actual cursor in the displayed image. This default zoom window cursor appears as a large "+" sign. The other dummy cursor that is available for the zoom window is smaller, and shaped like a box. Depending on the characteristics of the image that is currently displayed, one or the other of the available dummy cursors may be more appropriate. Alternatively, some users may not want the dummy" cursor to be displayed at all. The "Zoom Window Cursor" option is a cyclic option: the first click will change the dummy cursor from the large "+" to the small box; the second click will turn the dummy cursor off; and the third click will turn the dummy cursor back on, in the large "+" shape. After this, the cycle will repeat.

Update Mode

The Zooming function may be operated in two modes: Continuous or Button Press. When the update mode is Continuous, the zoom window is continually updated according to the cursor position in the displayed image. When the update mode is Button Press, the zoom window does not update until the mouse is clicked within the displayed image.

F. Print Pixel Values

This feature involves a workspace which displays the values of the pixels surrounding the pointer as the pointer is moved across the input image.

About The PrintPixel Display

The printpixel display will consist of a single (8x8) grid of boxes, or a set of three (8x8) grid of boxes, depending on whether or not an RGB image is displayed.

When the image displayed is *not* an RGB image, the numbers in the boxes represent the pixels surrounding the pointer at its current location in the image. Each box contains the numeric value of the pixel it represents; the background color of the box reflects the color of the pixel in the displayed image. The box in the grid that represents the actual pixel over which the pointer is positioned will be outlined so that it is easily spotted.

When an RGB image is displayed, the printpixel display contains one (8x8) grid for each of the R, G, and B planes in the image. The top grid displays the Red values; the middle grid represents the Green values; the bottom grid represents the blue values. The background color produced by each RGB value appears as the background of its corresponding box for all three of the red, green, and blue grids. The boxes in the three grids that represent the red, green, and blue values of the actual pixel over which the pointer is positioned are outlined so that they are easily recognized.

Update Mode

The update mode controls whether the printpixel display is updated on pointer motion or button press events. By default, the update mode is "Continuous"; ie, the pixel values in the print pixels display are updated continuously as the pointer is moved across the image. Changing the "Update mode:" logical selection to "Button Press" will cause the pixel values to be updated only when the button is pressed in the image.

G. Print Map Values

This subform contains a workspace which displays the values of the map data currently being displayed as red, green, and blue. The map values printed are those that are indexed by the pixels surrounding the pointer as the pointer is moved across the currently displayed image.

About the Print Map Values Display

The print map values workspace consists of three (8x8) grids of boxes. The first grid represents the values of the map column currently displayed as red, the second grid represents the values of the map column currently displayed as green, and the third grid represents the values of the map column currently displayed as blue. Note that you may change which map columns are currently displayed as red, green, and blue, by using the "Display" pane of the "Options" subform.

In each grid, the boxes represent the red, green, or blue values indexed by the values of the pixels surrounding the pointer at its current location in the image. Each box contains the numeric value of the red, green, or blue column that it represents; the background color of the box reflects the red, green, or blue element of the color of the pixel in the displayed image.

Update Mode

The update mode controls whether the print pixels display is updated on pointer motion or button press events. By default, the update mode is "Continuous"; ie, the pixel values in the print pixels display are updated continuously as the pointer is moved across the image. Changing the "Update mode:" logical selection to "Button Press" will cause the pixel values to be updated only when the button is pressed in the image.

H. Pseudo Color

The pseudocolor capability provides a method of interactively altering the colormap of the input image. First, you specify a single pixel or a range of pixels for which to alter the color; next, you specify the desired color.

About the Pseudocolor Display

The pseudocolor display includes a palette of the colors in which image pixels appear, and a set of Red, Green, and Blue scrollbars.

Pseudocolor Palette

The middle portion of the pseudocolor subform contains a palette, in the form of a grid of color cells. Each color cell represents a pixel in the input image, and appears in the same color as the pixel does in the image. The color cells are arranged in the order that the pixels appear in the image’s colormap; ie, the upper leftmost colorcell is the first color in the colormap and lower right-most pixel box is the last color in the colormap. When a range of colors from the palette is chosen for alteration, all colors that lie sequentially between those two color cells in the colormap are also altered.

The Red/Green/Blue Scrollbars

The three color bars below the palette will reflect the three different components that define the color of the selected pixel. These scroll bars are used to change the colormap of the input image.

First, select a single pixel value or a range of pixel values to have their color changed by clicking with the mouse on the color cells in the palette that display the pixel value(s) that you want to change. After selecting the pixels to be changed from the pseudocolor palette, use the scrollbars to specify the color in which that range of pixel values should appear.

I. Thresholding

The threshold feature of editimage provides a mechanism with which you may perform pixel thresholding or pixel clipping operations on the currently displayed image. When "Pixel Clipping" is the selected mode, you can perform pixel clipping on the image; when "Pixel Thresholding" is the selected mode, you can perform pixel thresholding on the image.

I.1. Pixel Clipping

Pixel clipping is used to mask out certain parts of the image so that only a desired range of pixels is observed. Pixel clipping works by allowing a windowed region of the original pixel values to be displayed. The user selects the pixel value of the masked region; all values outside the windowed region will be set to this masked region pixel value. If the masked region pixel value does not exist in the colormap, the corresponding grey level value will be used.

When the "Clipped Region" logical is set to "accept", the original pixel values of the image will be displayed *inside* the region defined by the lower and upper scroll bars; this can be thought of as a bandpass filter. When the "Clipped Region" logical is set to "reject", the original pixel values will be displayed *outside* the region defined by the lower and upper scroll bars; this can be thought of as a bandreject filter.

I.2. Pixel Thresholding

Pixel thresholding is used to perform real-time image thresholding. It generates a binary (or 2 pixel value) image. Pixel thresholding works by assigning a non-zero pixel value to the selected region; the rest of the image will be turned black.

When the "Threshold Region" logical is set to "non-inverted", the pixels in the image *inside* the region defined by the lower and upper scroll bars will be displayed as white, and those outside the region will be displayed as black. When the "Threshold Region" logical is set to "inverted", the opposite display is used, ie, those values *outside* the region defined by the lower and upper scroll bars will be displayed as white, and those inside the region will be displayed as black.

This page left intentionally blank

Table of Contents